

Accelerated Runtime Verification of LTL Specifications with Counting Semantics

Ramy Medhat, Yogi Joshi, Borzoo Bonakdarpour, and Sebastian Fischmeister
University of Waterloo, McMaster University, Canada
{rmedhat, y2joshi, sfischme}@uwaterloo.ca, borzoo@McMaster.ca

ABSTRACT

Runtime verification is an effective automated method for specification-based offline testing and analysis as well as on-line monitoring of complex systems. The specification language is often a variant of regular expressions or a popular temporal logic, such as LTL. This paper presents a novel and efficient parallel algorithm for verifying a more expressive version of LTL specifications that incorporates counting semantics, where nested quantifiers can be subject to numerical constraints. Such constraints are useful in evaluating thresholds (e.g., expected uptime of a web server). The significance of this extension is that it enables us to reason about the correctness of a large class of systems, such as web servers, OS kernels, and network behavior, where properties are required to be instantiated for parameterized requests, kernel objects, network nodes, etc. Our algorithm uses the popular *MapReduce* architecture to split a program trace into variable-based clusters at run time. Each cluster is then mapped to its respective monitor instances, verified, and reduced collectively on a multi-core CPU or the GPU. Our algorithm is fully implemented and we report very encouraging experimental results, where the monitoring overhead is negligible on real-world data sets.

1. INTRODUCTION

In this paper, we study runtime verification of properties specified in an extension of linear temporal logic (LTL) that supports expression of counting semantics with numerical constraints. Runtime verification (RV) is an automated specification-based technique, where a *monitor* evaluates the correctness of a set of logical properties on a particular execution either on the fly (i.e., at run time) or based on log files. Runtime verification complements exhaustive approaches such as model checking and theorem proving and under-approximated methods such as testing. The addition of counting semantics to properties is of particular interest, as they can express parametric requirements on types of execution entities (e.g., processes and threads), user- and kernel-level events and objects (e.g., locks, files, sockets), web services (e.g., requests and responses), and network traffic. For example, the requirement ‘every open file should eventually be closed’ specifies a rule for causal and temporal order of opening and closing individual objects which generalizes to *all* files. Such properties cannot be expressed using traditional RV frameworks, where the specification language is propositional LTL or regular expressions.

In this paper, we extend the 4-valued semantics of LTL (i.e., LTL_4), designed for runtime verification [3] by adding counting semantics with numerical constraints and propose an efficient parallel algorithm for their verification at run time. Inspired by the work in [13], the syntax of our language (denoted LTL_4-C) extends LTL syntax by the addition of

counting quantifiers. That is, we introduce two quantifiers: the *instance counting quantifier* (\mathbb{E}) which allows expressing properties that reason about the number of satisfied or violated instances, and the *percentage counting quantifier* (\mathbb{A}) which allows reasoning about the percentage of satisfied or violated instances out of all instances in a trace. These quantifiers are subscripted with numerical constraints to express the conditions used to evaluate the count. For example, the following LTL_4-C formula:

$$\mathbb{A}_{\geq 0.95} s : \text{socket}(s) \Rightarrow (\mathbf{G} \text{ receive}(s) \implies \mathbf{F} \text{ respond}(s))$$

intends to express the property that ‘at least 95% of open TCP/UDP sockets must eventually be closed’. also, the formula:

$$\mathbb{A}x : \text{user}(x) \Rightarrow (\mathbb{E}_{\leq 3} r : \text{rid}(r) \Rightarrow (\text{login} \wedge \text{unauthorized}))$$

intends to capture the requirement that ‘for all users, there exist at most 3 requests of type login that end with an unauthorized status’. The semantics of LTL_4-C is defined over six truth values:

- **True** (\top) denotes that the property is already permanently satisfied.
- **False** (\perp) denotes that the property is already permanently violated.
- **Currently true** (\top_c) denotes that the current execution satisfies the quantifier constraint of the property, yet it is possible that an extension violates the constraint.
- **Currently false** (\perp_c) denotes that the current execution violates the quantifier constraint of the property, yet it is possible that an extension satisfies it.
- **Presumably true** (\top_p) denotes that the current execution satisfies the inner LTL property and the quantifier constraint of the property.
- **Presumably false** (\perp_p) denotes that the current execution violates the inner LTL property and the quantifier constraint of the property.

We claim that these truth values provide us with informative verdicts about the status of different components of properties (i.e., quantifiers and their numerical constraints as well as the inner LTL formula) at run time.

The second contribution of this paper is a divide-and-conquer-based online monitor generation technique for LTL_4-C specifications. In fact, LTL_4-C monitors have to be generated at run time, otherwise, an enormous number of monitors (in the size of cross-product of domains of all

variables) has to be created statically, which is clearly impractical. Our technique first synthesizes an LTL₄ monitor for the inner LTL property of LTL₄-C properties pre-compile time using the technique in [3]. Then, based upon the values of variables observed at run time, submonitors are generated and merged to compute the current truth value of a property for the current program trace.

Our third contribution is an algorithm that implements the above approach for verification of LTL₄-C properties at run time. This algorithm enjoys two levels of parallelism: the monitor (1) works in parallel with the program under inspection, and (2) evaluates properties in a parallel fashion as well. While the former ensures that the runtime monitor does not intervene with the normal operation of the program under inspection, the latter attempts to maximize the throughput of the monitor. The algorithm utilizes the popular *MapReduce* technique to (1) spawn submonitors that aim at evaluating subformulas using partial quantifier elimination, and (2) merge partial evaluations to compute the current truth value of properties.

Our parallel algorithm for verification of LTL₄-C properties is fully implemented on multi-core CPU and GPU technologies. We report rigorous experimental results by conducting three real-world independent case studies. The first case study is concerned with monitoring HTTP requests and responses on an Apache Web Server. The second case study attempts to monitor users uploading maximum chunk packets repeatedly to a personal cloud storage service based on a dataset for profiling DropBox traffic. The third case study monitors a network proxy cache to reduce the bandwidth usage of online video services, based on a YouTube request dataset. We present performance results comparing single-core CPU, multi-core CPU, and GPU implementations. Our results show that our GPU-based implementation provides an average speed up of 7x when compared to single-core CPU, and 1.75x when compared to multi-core CPU. The CPU utilization of the GPU-based implementation is negligible compared to multi-core CPU, freeing up the system to perform more computation. Thus, the GPU-based implementation manages to provide competitive speedup while maintaining a low CPU utilization, which are two goals that the CPU cannot achieve at the same time. Put it another way, the GPU-based implementation incurs minimal monitoring costs while maintaining a high throughput.

The rest of the paper is organized as follows. Section 2 describes the syntax and semantics of LTL₄-C. In Section 3, we explain our online monitoring approach, while Section 4 presents our parallelization technique based on MapReduce. Experimental results are presented in Section 5. Related work is discussed in Section 6. Finally, we make concluding remarks and discuss future work in Section 7.

2. LTL WITH COUNTING SEMANTICS

To introduce our logic, we first define a set of basic concepts.

DEFINITION 1 (PREDICATE). Let $V = \{x_1, x_2, \dots, x_n\}$ be a set of variables with (possibly infinite) domains $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_n$, respectively. A predicate p is a binary-valued function on the domains of variables in V such that

$$p : \mathcal{D}_1 \times \mathcal{D}_2 \times \dots \times \mathcal{D}_n \rightarrow \{\text{true}, \text{false}\} \blacksquare$$

The arity of a predicate is the number of variables it accepts. A predicate is *uninterpreted* if the domain of variables are not known concrete sets. For instance, $p(x_1, x_2)$ is an uninterpreted predicate, yet we can interpret it as (for instance) a binary function that checks whether or not x_1 is less than x_2 over natural numbers.

Let UP be a finite set of uninterpreted predicates, and let $\Sigma = 2^{UP}$ be the power set of UP . We call each element of Σ an *event*.

DEFINITION 2 (TRACE). A trace $w = w_0 w_1 \dots$ is a finite or infinite sequence of events; i.e., $w_i \in \Sigma$, for all $i \geq 0$. ■

We denote the set of all infinite traces by Σ^ω and the set of all finite traces by Σ^* . A *program trace* is a sequence of events, where each event consists of *interpreted* predicates only. For instance, the following trace is a program trace:

$$w = \{\text{open}(1), r, \text{anony}\} \{\text{open}(2), rw, \text{user}(5)\} \dots$$

where *open* and *user* are unary predicates and *r*, *anony*, and *rw* are 0-arity predicates. Predicate *open* is interpreted as opening a file, *r* is interpreted as read-only permissions, *anony* is interpreted as an anonymous user, and so on.

2.1 Syntax of LTL₄-C

LTL₄-C extends LTL₄ with two counting quantifiers: the instance counting quantifier (\mathbb{E}) and the percentage counting quantifier (\mathbb{A}). The semantics of these quantifiers are introduced in subsection 2.4. The syntax of LTL₄-C is defined as follows:

DEFINITION 3 (LTL₄-C SYNTAX). LTL₄-C formulas are defined using the following grammar:

$$\begin{aligned} \varphi &::= \mathbb{A}_{\sim k} x : p(x) \Rightarrow \varphi \mid \mathbb{E}_{\sim l} x : p(x) \Rightarrow \varphi \mid \psi \\ \psi &::= \top \mid p(x_1 \dots x_n) \mid \neg \psi \mid \psi_1 \wedge \psi_2 \mid \\ &\quad \mathbf{X} \psi \mid \psi_1 \mathbf{U} \psi_2 \end{aligned}$$

where \mathbb{A} is the percentage counting quantifier, \mathbb{E} is the instance counting quantifier, $x, x_1 \dots x_n$ are variables with possibly infinite domains $\mathcal{D}, \mathcal{D}_1, \dots, \mathcal{D}_n$, $\sim \in \{<, \leq, >, \geq, =\}$, $k : \mathbb{R} \in [0, 1]$, $l \in \mathbb{Z}^+$, \mathbf{X} is the next, and \mathbf{U} is the until temporal operators. ■

If we omit the numerical constraint in $\mathbb{A}_{\sim k}$ (respectively, $\mathbb{E}_{\sim l}$), we mean $\mathbb{A}_{=1}$ (respectively, $\mathbb{E}_{\geq 1}$). The syntax of LTL₄-C forces constructing formulas, where a string of counting quantifiers is followed by a quantifier-free formula. We emphasize that \mathbb{A} and \mathbb{E} do not necessarily resemble standard first-order quantifiers \forall and \exists . In fact, as we will explain $\neg \mathbb{A}$ and \mathbb{E} are not generally equivalent.

Consider LTL₄-C property $\varphi = \mathbb{A} x : p(x) \Rightarrow \psi$, where the domain of x is \mathcal{D} . This property denotes that for any possible valuation of the variable x ($[x := v]$), if $p(v)$ holds, then ψ should hold. If $p(v)$ does not hold, then $p(v) \Rightarrow \psi$ trivially evaluates to true. This effectively means that the quantifier $\mathbb{A} x$ is in fact applied only over the following sub-domain:

$$\{v \in \mathcal{D} \mid p(v)\} \subseteq \mathcal{D}$$

To give an intuition, consider the scenarios where file management anomalies can cause serious problems at run time (e.g., in NASA's Spirit Rover on Mars in 2004). For example, the following LTL₄-C property expresses "at least half of the files that a process has previously opened must be closed":

$$\varphi_1 = \mathbb{A}_{\geq 50\%} f : \text{intrace}(f) \Rightarrow (\text{opened}(f) \mathbf{U} \text{close}(f)) \quad (1)$$

where *intrace* denotes the fact that the concrete file appeared in any event in the trace.

2.2 4-Valued LTL [3]

First, we note that the syntax of LTL₄ can be easily obtained from Definition 3 by (1) removing the counting quantifier rules and (2) reducing the arity of predicates to 0 (i.e., predicates become atomic propositions).

2.2.1 FLTL

To introduce LTL_4 semantics, we first introduce Finite LTL. Finite LTL (FLTL) [14] allows us to reason about finite traces for verifying properties at run time. The semantics of FLTL is based on the truth values $\mathbb{B}_2 = \{\top, \perp\}$.

DEFINITION 4 (FLTL SEMANTICS). Let φ and ψ be LTL properties, and $u = u_0 u_1 \dots u_{n-1}$ be a finite trace.

$$[u \models_F \mathbf{X} \varphi] = \begin{cases} [u_1 \models_F \varphi] & \text{if } u_1 \neq \epsilon \\ \perp & \text{otherwise} \end{cases}$$

$$[u \models_F \varphi \mathbf{U} \psi] = \begin{cases} \top & \exists k \in [0, n-1] : [u_k \models_F \psi] = \top \wedge \\ & \forall l \in [0, k] : [u^l \models_F \varphi] = \top \\ \perp & \text{otherwise} \end{cases}$$

where ϵ is the empty trace. The semantics of FLTL for atomic propositions and Boolean combinations are identical to that of LTL. ■

Similar to standard LTL, $\mathbf{F}p \equiv \top \mathbf{U} p$ and $\mathbf{G}p \equiv \neg \mathbf{F} \neg p$.

2.2.2 LTL₄ Semantics

LTL_4 is designed for runtime verification by producing more informative verdicts than FLTL. The semantics of LTL_4 is defined based on values $\mathbb{B}_4 = \{\top, \top_p, \perp_p, \perp\}$ (*true*, *presumably true*, *presumably false*, and *false* respectively). The semantics of LTL_4 is defined based on the semantics LTL and FLTL.

DEFINITION 5 (LTL₄ SEMANTICS). Let φ be an LTL_4 property and u be a finite prefix of a trace.

$$[u \models_4 \varphi] = \begin{cases} \top & \forall v \in \Sigma^\omega : uv \models \varphi \\ \perp & \forall v \in \Sigma^\omega : uv \not\models \varphi \\ \top_p & [u \models_F \varphi] \wedge \exists v \in \Sigma^\omega : uv \not\models \varphi \\ \perp_p & [u \not\models_F \varphi] \wedge \exists v \in \Sigma^\omega : uv \models \varphi \end{cases} \blacksquare$$

In this definition, \models denotes the satisfaction relation defined by standard LTL semantics over infinite traces. Thus, an LTL_4 property evaluates to \top with respect to a finite trace u , if the property remains *permanently satisfied*, meaning that for all possible infinite continuations of the trace, the property will always be satisfied in LTL. Likewise, a valuation of \perp means that the property will be *permanently violated*. If the property evaluates to \top_p , this denotes that currently the property is satisfied yet there exists a continuation that could violate it. Finally, value \perp_p denotes that currently the property is violated yet there exists a continuation that could satisfy it.

2.2.3 LTL₄ Monitors

In [3], the authors introduce a method of synthesizing a *monitor*, as a deterministic finite state machine (FSM), for an LTL_4 property.

DEFINITION 6 (LTL₄ MONITOR). Let φ be an LTL_4 formula over Σ . The monitor \mathcal{M}_φ of φ is the unique FSM $(\Sigma, Q, q_0, \delta, \lambda)$, where Q is a set of states, q_0 is the initial state, $\delta \subseteq Q \times \Sigma \times Q$ is the transition relation, and $\lambda : Q \rightarrow \mathbb{B}_4$ is a function such that:

$$[u \models_4 \varphi] = \lambda(\delta(q_0, u)). \blacksquare$$

Thus, given an LTL_4 property φ and a finite trace u , monitor \mathcal{M}_φ is capable of producing a truth value in \mathbb{B}_4 , which is equal to $[u \models_4 \varphi]$. For example, Figure 1 shows the monitor for property $\varphi = \mathbf{G}a \vee (b \mathbf{U} c)$. Observe that a monitor




Figure 1: LTL_4 monitor for property $\varphi = \mathbf{G}a \vee (b \mathbf{U} c)$.

has two *trap* states (only an outgoing self loop), which map to truth values \top and \perp . They are trap states since these truth values imply permanent satisfaction (respectively, violation). Otherwise, states labeled by \top_p and \perp_p can have outgoing transitions to other states.

2.3 Truth Values of LTL₄-C

The objective of $LTL_4\text{-C}$ is to verify the correctness of quantified properties at run time with respect to finite program traces. Such verification attempts to produce a sound verdict regardless of future continuations.

We incorporate six truth values to define the semantics of $LTL_4\text{-C}$: $\mathbb{B}_6 = \{\top, \perp, \top_c, \perp_c, \top_p, \perp_p\}$; *true*, *false*, *currently true*, *currently false*, *presumably true*, *presumably false*, respectively. The values in \mathbb{B}_6 form a lattice ordered as follows: $\perp < \perp_c < \perp_p < \top_p < \top_c < \top$. Given a finite trace u and an $LTL_4\text{-C}$ property φ , the informal description of evaluation of u with respect to φ is as follows:

- **True** (\top) denotes that any infinite extension of u satisfies φ .
- **False** (\perp) denotes that any infinite extension of u violates φ .
- **Currently true** (\top_c) denotes that currently u satisfies the counting quantifier constraint of φ , yet it is possible that a suffix of u violates the constraint. For instance, the valuation of Property 1 (i.e., φ_1) is \top_c , if in a trace u , currently 50% of files previously opened are closed. This is because (1) the inner LTL property is permanently satisfied for at least 50% of files previously opened, and (2) it is possible for a trace continuation to change this percentage to less than 50% in the future (a trace in which enough new files are opened and not closed).
- **Currently false** (\perp_c) denotes that currently u violates the quantifier constraint of φ , yet it is possible that a suffix of u satisfies the constraint. For instance, the valuation of Property 1 (i.e., φ_1) in a finite trace u is \perp_c , if the number of files that were not successfully opened is currently greater than 50%. This could happen in the scenario where opening a file fails, possibly due to lack of permissions. Analogous to \top_c , the property is evaluated to \perp_c because (1) the inner LTL property is permanently satisfied for less than 50% of files in the program trace, and (2) it is possible for a trace continuation to change this percentage to at least 50% in the future.

Now let us consider modifying the property to support multiple open and close operations on the same file. For this purpose, we reformulate the property as follows:

$$\varphi_2 = \mathbb{A}_{\geq 50\%} f : \text{intrace}(f) \Rightarrow (\mathbf{G}(\text{opened}(f) \mathbf{U} \text{close}(f))) \quad (2)$$

- **Presumably true** (\top_p) extends the definition of *presumably true* in LTL_4 [4], where \top_p denotes that u sat-

ifies the inner LTL property and the counting quantifier constraint in φ , if the program terminates after execution of u . For example, Property 2 (i.e., φ_2) evaluates to \top_p , if at least 50% of the files in the program trace are closed. Closed files presumably satisfy the property, since they satisfy the **G** operator thus far, yet can potentially violate it if the file is opened a subsequent time without being closed. Note that this property can never evaluate to \top_c , since no finite trace prefix can permanently satisfy the inner LTL property. However, if the inner property can be permanently satisfied (\top) and presumably satisfied (\top_p), then the entire $\text{LTL}_4\text{-C}$ property can potentially evaluate to \top_c if the numerical condition of the quantifier is satisfied. A property can evaluate to \top_p only if the conditions for \top_c are not met, since \top_c is higher up the partial order of \mathbb{B}_6 .

- **Presumably false** (\perp_p) extends the definition of *presumably false* in LTL_4 [4], which denotes that u presumably violates the quantifier constraint in φ . According to the Property 2, this scenario will occur when the number of files that are either closed or opened and not yet closed is at least 50% of all files in the trace. Opened files presumably violate the inner property, since closing the file is required but has not yet occurred. This condition should not conflict with \top_p or \top_c , since they precede \perp_p in the partial order of \mathbb{B}_6 and thus \perp_p only occurs if the conditions for \top_p and \top_c do not hold.

2.4 Semantics of $\text{LTL}_4\text{-C}$

An $\text{LTL}_4\text{-C}$ property essentially defines a set of traces, where each trace is a sequence of events (i.e., sets of uninterpreted predicates). We define the semantics of $\text{LTL}_4\text{-C}$ with respect to finite traces and present a method of utilizing these semantics for runtime verification. In the context of runtime verification, the objective is to ensure that a program trace (i.e., a sequence of sets of *interpreted* predicates) is in the set of traces that the property defines, given the interpretations of the property predicates within the program trace.

To introduce the semantics of $\text{LTL}_4\text{-C}$, we examine counting quantifiers further. Since the syntax of $\text{LTL}_4\text{-C}$ allows nesting of counting quantifiers, a canonical form of properties is as follows:

$$\varphi = \mathbb{Q}_\varphi \psi \quad (3)$$

where ψ is an LTL property and \mathbb{Q}_φ is a string of counting quantifiers

$$\mathbb{Q}_\varphi = \mathcal{Q}_0 \mathcal{Q}_1 \cdots \mathcal{Q}_{n-1} \quad (4)$$

such that each $\mathcal{Q}_i = \langle \mathcal{Q}_i, \sim_i, c_i, x_i, p_i \rangle$, $0 \leq i \leq n-1$, is a tuple encapsulating the counting quantifier information. That is, $\mathcal{Q}_i \in \{\mathbb{A}, \mathbb{E}\}$, $\sim_i \in \{<, \leq, >, \geq, =\}$, c_i is the constraint constant, x_i is the bound variable, and p_i is the predicate within the quantifier (see Definition 3).

We present semantics of $\text{LTL}_4\text{-C}$ in a stepwise manner:

1. **Variable valuation.** First, we demonstrate how variable valuations are extracted from the trace and used to substitute variables in the formula.
2. **Canonical variable valuations.** Next, we demonstrate how to build a canonical structure of the variable valuations provided in Step 1. This canonical structure mirrors the canonical structure of $\text{LTL}_4\text{-C}$ properties.

3. **Valuation of property instances.** A *property instance* is a unique substitution of variables in the property with values from their domains. This step demonstrates how to evaluate property instances.

4. **Applying quantifier numerical constraints.** This step demonstrates how to evaluate counting quantifiers by applying their numerical constraints on the valuation of a set of property instances from Step 3. The set of property instances is retrieved with respect to the canonical structure defined in Step 2.

5. **Inductive semantics.** Using the canonical structure in Step 2, and valuation of counting quantifiers in Step 4, we define semantics that begin at the outermost counting quantifier of an $\text{LTL}_4\text{-C}$ property and evaluate quantifiers recursively inwards.

2.4.1 Variable Valuation

We define a vector D_φ with respect to a property φ as follows:

$$D_\varphi = \langle d_0, d_1, \dots, d_{n-1} \rangle$$

where $n = |\mathbb{Q}_\varphi|$ and d_i , $0 \leq i \leq n-1$, is a value for variable x_i . We denote the first m components of the vector D_φ (i.e., $\langle d_0, d_1, \dots, d_{m-1} \rangle$) by $D_\varphi|_m^m$. We refer to D_φ as a *value vector* and to $D_\varphi|_m^m$ as a *partial value vector*.

A *property instances* $\hat{\varphi}(D_\varphi|_m^m)$ is obtained by replacing every occurrence of the variables $x_0 \cdots x_{m-1}$ in φ with the values $d_0 \cdots d_{m-1}$, respectively. Thus, $\hat{\varphi}(D_\varphi|_m^m)$ is free of quantifiers of index less than m , yet remains quantified over variables $x_m \cdots x_{n-1}$. For instance, for the following property

$$\varphi = \mathbb{A}_{>c_1} x : p_x(x) \Rightarrow (\mathbb{A}_{<c_2} y : p_y(y) \Rightarrow \mathbf{G} q(x, y))$$

and value vector $D_\varphi = \langle 1, 2 \rangle$ (i.e., the vector of values for variables x and y , respectively), $\hat{\varphi}(D_\varphi)$ will be

$$\hat{\varphi}(\langle 1, 2 \rangle) = p_x(1) \Rightarrow (p_y(2) \Rightarrow \mathbf{G} q(1, 2))$$

We now define the set $\mathbb{D}_{\varphi, u}$ as the set of all value vectors with respect to a property $\varphi = \mathbb{Q}_\varphi \psi$ and a finite trace $u = u_0 u_1 \cdots u_k$:

$$\mathbb{D}_{\varphi, u} = \{D_\varphi \mid \exists j \in [0, k] : \forall i \in [0, n-1] : p_i(d_i) \in u_j\} \quad (5)$$

where $n = |\mathbb{Q}_\varphi|$.

2.4.2 Canonical Variable Valuations

An $\text{LTL}_4\text{-C}$ property follows a canonical structure, in which every counting quantifier \mathcal{Q}_i has a *parent* quantifier \mathcal{Q}_{i-1} , except for \mathcal{Q}_0 which is the *root* counting quantifier. A counting quantifier \mathcal{Q}_i is applied over all valuations of its variable x_i given a unique valuation of its predecessor variables x_0, \dots, x_{i-1} . Hence, we define function \mathcal{P} which takes as input a partial value vector $D_\varphi|_m^m$, and returns all partial value vectors in $\mathbb{D}_{\varphi, u}$ of length $m+1$, such that the first m elements of these vectors is the same as $D_\varphi|_m^m$. In this context, we refer to $D_\varphi|_m^m$ as a *parent* vector and all the returned vectors as *child* vectors. Similarly, a property instance can have a parent; for instance, $\hat{\varphi}(D_\varphi|_m^m)$ is the parent of $\hat{\varphi}(D_\varphi|_{m+1}^{m+1})$.

$$\mathcal{P}(\varphi, u, D_\varphi|_m^m) = \left\{ D'_\varphi|_{m+1}^{m+1} \mid D'_\varphi \in \mathbb{D}_{\varphi, u} \wedge D'_\varphi|_m^m = D_\varphi|_m^m \right\}$$

Following the example above, assume there are two value vectors: $\langle 1, 2 \rangle$ and $\langle 1, 3 \rangle$. In this case,

$$\mathcal{P}(\varphi, u, \langle 1 \rangle) = \{ \langle 1, 2 \rangle, \langle 1, 3 \rangle \}$$

2.4.3 Valuation of Property Instances

As per the definition of $\mathbb{D}_{\varphi, u}$, every value vector $D_\varphi = \langle d_0 \dots d_{n-1} \rangle$ in $\mathbb{D}_{\varphi, u}$ contains values for which the predicates $p_i(d_i)$ hold in some trace event u_j . For simplicity, we denote this as a value vector *in* a trace event u_j . These value vectors can possibly be in multiple and interleaved events in the trace. Thus, we define a trace $u^{D_\varphi} = u_0^{D_\varphi} u_1^{D_\varphi} \dots u_l^{D_\varphi}$ as a subsequence of the trace u such that the value vector D_φ is in every event:

$$\forall j \in [0, l] : \forall i \in [0, n-1] : p_i(d_i) \in u_j^{D_\varphi}$$

For any property instance $\hat{\varphi}(D_\varphi)$, we wish to evaluate $[u^{D_\varphi} \models_6 \hat{\varphi}(D_\varphi)]$ (read as valuation of $\hat{\varphi}(D_\varphi)$ with respect to u^{D_φ} for LTL₄-C), since any other event in trace u is not of interest to $\hat{\varphi}(D_\varphi)$.

By leveraging u^{D_φ} , we define function \mathcal{B} as follows:

$$\mathcal{B}(\varphi, u, D_\varphi |^m, b) = \begin{cases} D_\varphi' |^{m+1} \in \mathcal{P}(\varphi, u, D_\varphi |^m) \mid \\ [u^{D_\varphi' |^{m+1}} \models_6 \hat{\varphi}(D_\varphi' |^{m+1})] = b & \text{iff } m < |\mathbb{Q}_\varphi| - 1 \\ D_\varphi' |^{m+1} \in \mathcal{P}(\varphi, u, D_\varphi |^m) \mid \\ [u^{D_\varphi' |^{m+1}} \models_4 \hat{\varphi}(D_\varphi' |^{m+1})] = b & \text{iff } m = |\mathbb{Q}_\varphi| - 1 \end{cases}$$

where b is a truth value in \mathbb{B}_6 . Function \mathcal{B} can be implemented in a straightforward manner, where it iterates over all its children value vectors $D_\varphi' |^{m+1}$ which are retrieved using \mathcal{P} . For every child vector, the function checks whether $\hat{\varphi}(D_\varphi' |^{m+1})$ evaluates to b with respect to the trace subsequence $u^{D_\varphi' |^{m+1}}$.

To clarify \mathcal{B} , let us refer to our example earlier. Let a program trace u be as follows:

$$u = \{p_x(1), p_y(2), \dots\}, \{p_x(1), p_y(3), \dots\}, \{p_x(1), p_y(2), \dots\}$$

With respect to this trace, $\mathcal{P}(\varphi, u, \langle 1 \rangle) = \{\langle 1, 2 \rangle, \langle 1, 3 \rangle\}$. As per the definition of u^{D_φ} , $u^{(1,2)} = u_0 u_2$, and $u^{(1,3)} = u_1$. Thus, $\mathcal{B}(\varphi, u, \langle 1 \rangle, b)$ checks the following:

$$\begin{aligned} [u^{(1,2)} \models_4 p_x(1) \Rightarrow (p_y(2) \Rightarrow \mathbf{G} q(1, 2))] &= b \\ [u^{(1,3)} \models_4 p_x(1) \Rightarrow (p_y(3) \Rightarrow \mathbf{G} q(1, 3))] &= b \end{aligned}$$

The definition of u^{D_φ} implies that $p_i(d_i) \in u_j^{D_\varphi}$ for all j . Thus, we can simplify the property by omitting the p predicates since they hold by definition:

$$\begin{aligned} [u^{(1,2)} \models_4 \mathbf{G} q(1, 2)] &= b \\ [u^{(1,3)} \models_4 \mathbf{G} q(1, 3)] &= b \end{aligned}$$

For instance, if only $[u^{(1,2)} \models_4 \mathbf{G} q(1, 2)] = b$ holds, then

$$\mathcal{B}(\varphi, u, \langle 1 \rangle, b) = \{\langle 1, 2 \rangle\}$$

As can be seen in the example, the property instances that are evaluated are LTL₄ properties. This is because the input to \mathcal{B} is $D_\varphi |^1 = D_\varphi |^{|\mathbb{Q}_\varphi|-1}$, which represents the inner most quantifier.

2.4.4 Applying Quantifier Numerical Constraints

Finally, numerical constraints should be incorporated in the semantics. We define function \mathcal{S} as follows:

$$\mathcal{S}(\varphi, u, D_\varphi |^m, B) = \begin{cases} \bigcup_{b \in B} \mathcal{B}(\varphi, u, D_\varphi |^m, b) \mid c_i \times |\{\mathcal{P}(\varphi, u, D_\varphi |^m)\}| & \text{iff } Q_m = \mathbb{A} \\ \bigcup_{b \in B} \mathcal{B}(\varphi, u, D_\varphi |^m, b) \mid \sim_i c_i & \text{iff } Q_m = \mathbb{E} \end{cases} \quad (6)$$

where $B \subseteq \mathbb{B}_6$ is a set of truth values. This function returns whether a counting quantifier constraint is satisfied or not based on any of the truth values $b \in B$. Observe that, for percentage counting quantifiers, the constraint value denotes the percentage of property instances that evaluate to b . For instance counting quantifiers, the constraint value denotes the number of property instances that evaluate to b . For instance, consider Property 7 which is read as: for all users, there exists at most 3 requests of type login that end with an unauthorized status. For such a property, if 4 or more unauthorized login attempts are detected for the same user, the property is permanently violated.

2.4.5 Inductive Semantics

Using the previously defined set of functions, we now formalize LTL₄-C semantics.

DEFINITION 7 (LTL₄-C SEMANTICS). LTL₄-C semantics for properties with counting quantifiers are defined as follows:

$$[u \models \varphi] = \begin{cases} \top & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \{\top\}) = 1 \wedge \\ & \forall v \in \Sigma^\omega : [uv \models_6 \varphi] = \top \\ \perp & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \mathbb{B}_6 - \{\perp\}) = 0 \wedge \\ & \forall v \in \Sigma^\omega : [uv \models_6 \varphi] = \perp \\ \top_c & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \{\top, \top_c\}) = 1 \wedge \\ & \exists v \in \Sigma^\omega : [uv \models_6 \varphi] \neq \top_c \\ \perp_c & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \mathbb{B}_6 - \{\perp, \perp_c\}) = 0 \wedge \\ & \exists v \in \Sigma^\omega : [uv \models_6 \varphi] \neq \perp_c \\ \top_p & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \{\top, \top_c, \top_p\}) = 1 \wedge \\ & \mathcal{S}(\varphi, u, \langle \rangle, \{\top, \top_c\}) = 0 \\ \perp_p & \text{iff } \mathcal{S}(\varphi, u, \langle \rangle, \{\top, \top_c, \top_p\}) = 0 \wedge \\ & \mathcal{S}(\varphi, u, \langle \rangle, \mathbb{B}_6 - \{\perp, \perp_c\}) = 0 \blacksquare \end{cases}$$

Note that these semantics are applied recursively until there is only one counting quantifier left in the formula, at which point \mathcal{B} checks the valuation based on LTL₄ semantics ($[u^{D_\varphi} \models_4 \hat{\varphi}(D_\varphi)] = b$). When checking the valuation of these LTL₄ properties, \mathcal{B} will always return an empty set in case the input b is \top_c or \perp_c , since these truth values are inapplicable to LTL₄ properties. As mentioned earlier, truth values in \mathbb{B}_6 form a lattice. Standard lattice operators \sqcap and \sqcup are defined as expected based on the lattice's partial order. Permanent satisfaction (\top) or violation (\perp) is applicable to \mathbb{E} quantifiers regardless of the comparison operator, as well as a special case of \mathbb{A} quantifiers:

- **\mathbb{A} quantifier.** As mentioned earlier, if the \mathbb{A} quantifier is not subscripted, it is assumed to denote $\mathbb{A}_{=1}$. In this case, a single violation in its child property instances causes a permanent violation of the quantified property.
- **\mathbb{E} quantifier.** Permanent violation is possible for any numerical constraint attached to an \mathbb{E} quantifier, since it is a condition on the *number* of satisfied property instances.

Property 7 illustrates an example of an \mathbb{E} quantifier that can be permanently violated. Also, since the \mathbb{A} quantifier in Property 7 defaults to $\mathbb{A}_{=1}$, it will be violated if a single user makes more than three unauthorized login attempts. In such a case, the entire property evaluates to \perp . Table 1 illustrates how permanent satisfaction or violation apply to the different numerical constraints of \mathbb{E} quantifiers.

$$\mathbb{A}x : \text{user}(x) \Rightarrow (\mathbb{E}_{\leq 3} r : \text{rid}(r) \Rightarrow (\text{login} \wedge \text{unauthorized})) \quad (7)$$

Table 1: Rules of permanent satisfaction or violation of \mathbb{E} constraints

Operator	Verdict
$> c$	Permanent satisfaction if $> c$
$\geq c$	Permanent satisfaction if $\geq c$
$= c$	Permanent violation if $> c$
$< c$	Permanent violation if $\geq c$
$\leq c$	Permanent violation if $> c$

To clarify the semantics, consider Property 7 and the following program trace:

{rid(12), user(Adam), login, unauthorized}
 {rid(13), user(Adam), login, unauthorized}
 {rid(14), user(Jack), login, authorized}
 {rid(15), user(Adam), login, unauthorized}
 {rid(16), user(Adam), login, unauthorized}

where each line represents an event: a set of interpreted predicates. Each event contains a request identifier (rid), a username, a request type (login), and response status (authorized or unauthorized). As seen in the trace, there are 5 distinct value vectors: $\langle Adam, 12 \rangle$, $\langle Adam, 13 \rangle$, $\langle Jack, 14 \rangle$, $\langle Adam, 15 \rangle$, and $\langle Adam, 16 \rangle$. Now, let us apply the inductive semantics on the property.

Step 1. We begin by checking the truth value of $[u \models_6 \varphi]$, which requires determining which condition in Definition 7 applies. This requires the evaluation of function \mathcal{S} for the different truth values shown. Since we are verifying φ , we begin with the outermost counting quantifier, which is a \mathbb{A} quantifier. Thus, \mathcal{S} will require calculating the cardinality of the set $\mathcal{P}(\varphi, u, D_\varphi|_0)$, which in case of the trace should be $|\{Adam, Jack\}| = 2$. Now, in order to evaluate \mathcal{S} , one has to evaluate \mathcal{B} to determine whether each property instance evaluates to a certain truth value or not. The two property instances thus far are:

$$\hat{\varphi}(D_\varphi|_1) = \hat{\varphi}(Adam) = \mathbb{E}_{\leq 3} r : \text{rid}(r) \Rightarrow (\text{login} \wedge \text{unauthorized})$$

$$\hat{\varphi}(D'_\varphi|_1) = \hat{\varphi}(Jack) = \mathbb{E}_{\leq 3} r : \text{rid}(r) \Rightarrow (\text{login} \wedge \text{unauthorized})$$

And the trace subsequences for these property instances respectively are:

$$u^{D_\varphi|_1} = \{\text{rid}(12), \dots\} \{\text{rid}(13), \dots\} \{\text{rid}(15), \dots\} \{\text{rid}(16), \dots\}$$

$$u^{D'_\varphi|_1} = \{\text{rid}(14), \dots\}$$

Note that $\text{user}(Adam) \Rightarrow \dots$ is omitted from $\hat{\varphi}(D_\varphi|_1)$ since $\text{user}(Adam)$ holds according to the trace subsequence. The same applies to $\text{user}(Jack)$. Evaluating these property instances with respect to the trace subsequences requires referring to Definition 7 again, which marks the second level of recursion.

Step 2. Let us consider the property instance $\hat{\varphi}(D_\varphi|_1)$, which begins with an \mathbb{E} quantifier and has LTL₄ properties as child instances (refer to \mathcal{P}). These properties are in the form of $\text{login} \wedge \text{unauthorized}$, where there is one instance for each distinct request identifier. We can deduce that the property holds for all 4 requests: 12, 13, 15, and 16, thus evaluating to \top . Therefore, the following holds:

$$\mathcal{B}(\hat{\varphi}(D_\varphi|_1), u^{D_\varphi|_1}, D_\varphi|_1, \top) = 4$$

This value, when used in $\mathcal{S}(\hat{\varphi}(D_\varphi|_1), u^{D_\varphi|_1}, D_\varphi|_1, \{\top\})$ will violate the numerical condition: $4 \not\leq 3$, resulting in \mathcal{S} evaluating to 0 (false). Based on the conditions in Definition 7 and the rules of permanent violation, this property instance becomes permanently violated and thus the verdict is \perp .

The other property instance $\hat{\varphi}(D'_\varphi|_1)$ will however evaluate to \top since its child property instance

$$\hat{\varphi}(D'_\varphi|_2) = \hat{\varphi}(\langle Jack, 14 \rangle) = \text{login} \wedge \text{unauthorized}$$

is violated, and thus the number of satisfied instances is still less than 3.

Step 3. In this step we use the valuations determined in Step 2 to produce verdicts for the property instances in Step 1. Based on \mathcal{S} , the \mathbb{A} quantifier's numerical condition is violated, since not *all* instances are satisfied. The final verdict should thus be $[u \models_6 \varphi] = \perp$, which denotes a permanent violation of the property.

3. DIVIDE-AND-CONQUER-BASED MONITORING OF LTL₄-C

In this section, we describe our technique inspired by divide-and-conquer for evaluating LTL₄-C properties at run time. This approach forms the basis of our parallel verification algorithm in Section 4.

Unlike runtime verification of propositional LTL₄ properties, where the structure of a monitor is determined solely based on the property itself, a monitor for an LTL₄-C needs to evolve at run time, since the valuation of quantified variables change over time. More specifically, the monitor \mathcal{M}_φ for an LTL₄-C property $\varphi = \mathbb{Q}_\varphi \psi$ relies on instantiating a *submonitor* for each property instance $\hat{\varphi}$ obtained at run time. We incorporate two type of submonitors: (1) LTL₄ *submonitors* evaluate the inner LTL property ψ , and (2) *quantifier submonitors* deal with quantifiers in \mathbb{Q}_φ , described in Subsections 3.1 and 3.2. In Subsection 3.3, we explain the conditions under which a submonitor is instantiated at run time. Finally, in Subsection 3.4, we elaborate on how submonitors evaluate an LTL₄-C property.

3.1 LTL₄ Submonitors

Let $\varphi = \mathbb{Q}_\varphi \psi$ be an LTL₄-C property. If $|\mathbb{Q}_\varphi| = 0$ (respectively, one wants to evaluate $\hat{\varphi}(D_\varphi|_i)$, where $i = |\mathbb{Q}_\varphi|$), then φ (respectively, $\hat{\varphi}(D_\varphi|_i)$) is free of quantifiers and, thus, the monitor (respectively, submonitor) of such a property is a standard LTL₄ monitor (see Definition 6). We denote LTL₄ submonitors as $\mathcal{M}_{D_\varphi}^*$, where D_φ is the value vector with which the monitor is initialized.

3.2 Quantifier Submonitors

Given a finite trace u and an LTL₄-C property $\varphi = \mathbb{Q}_\varphi \psi$, a *quantifier submonitor* ($\mathcal{M}^\mathbb{Q}$) is a monitor responsible for determining the valuation of a property instance $\hat{\varphi}(D_\varphi|_i)$ with respect to a trace subsequence $u^{D_\varphi|_i}$, if $i < |\mathbb{Q}_\varphi|$. Obviously, such a valuation is in \mathbb{B}_6 . Let \mathbb{V} be a six-dimensional vector space, where each dimension represents a truth value in \mathbb{B}_6 .

DEFINITION 8 (QUANTIFIER SUBMONITOR). *Let $\varphi = \mathbb{Q}_\varphi \psi$ be an LTL₄-C property and $\hat{\varphi}(D_\varphi|_i)$ be a property instance, with $i \in [0, |\mathbb{Q}_\varphi| - 1]$. The quantifier submonitor for $\hat{\varphi}(D_\varphi|_i)$ is the tuple $\mathcal{M}_{D_\varphi|_i}^\mathbb{Q} = \langle \mathcal{Q}_i, \mathbb{M}_{D_\varphi|_i}, v, b \rangle$, where*

- \mathcal{Q}_i encapsulates the quantifier information (see Equation 4)
- $v \in \mathbb{V}$ represents the current number of child property instances that evaluate to each truth value in \mathbb{B}_6 with respect to their trace subsequences,
- $b \in \mathbb{B}_6$ is the current value of $[u^{D_\varphi|_i} \models_6 \hat{\varphi}(D_\varphi|_i)]$,
- $\mathbb{M}_{D_\varphi|_i}$ is the set of child submonitors (submonitors of

child property instances) defined as follows:

$$\mathbb{M}_{D_\varphi|^i} = \begin{cases} \{\mathcal{M}_{D_\varphi'}^* \mid D_\varphi'^i = D_\varphi|^i\} & \text{if } i = |\mathbb{Q}_\varphi| - 1 \\ \{\mathcal{M}_{D_\varphi'^{i+1}}^\varnothing \mid D_\varphi'^{i+1} = D_\varphi|^i\} & \text{if } i < |\mathbb{Q}_\varphi| - 1 \end{cases}$$

Thus, if $i = |\mathbb{Q}_\varphi| - 1$, all child submonitors are LTL_4 submonitors. Otherwise, they are quantifier submonitors of the respective child property instances. ■

Based on the definition, every quantifier submonitor references a set of child monitors. We use the following notation to denote a hierarchy of a submonitor and its children:

$$\mathcal{M}_{D_\varphi|^i}^\varnothing \{ \mathcal{M}_{D_\varphi'^{i+1}}^\varnothing, \mathcal{M}_{D_\varphi''^{i+1}}^\varnothing, \mathcal{M}_{D_\varphi'''^{i+1}}^\varnothing, \dots \}$$

such that $D_\varphi|^i = D_\varphi'^i = D_\varphi''^i \dots$ and $i < |\mathbb{Q}_\varphi| - 1$ which is why the child monitors are quantifier submonitors.

3.3 Instantiating Submonitors

Let an $\text{LTL}_4\text{-C}$ monitor \mathcal{M}_φ for property φ evaluate the property with respect to a finite trace $u = u_0u_1\dots$. Let $D_\varphi = \langle d_0, d_1, \dots \rangle$ be a value vector and u_k the first trace event such that $\forall d_i : p_i(d_i) \in u_k$, where p_i is the predicate within each quantifier (i.e. $\mathbb{A}x_i : p_i(x_i) \Rightarrow \dots$). In this case, the $\text{LTL}_4\text{-C}$ monitor instantiates submonitors for every property instance resulting from that value vector. A value vector of length $|\mathbb{Q}_\varphi|$ results in $|\mathbb{Q}_\varphi| + 1$ property instances: one for each quantifier in addition to an LTL_4 inner property. The hierarchy of the instantiated submonitors is as follows:

$$\mathcal{M}_{D_\varphi|^0}^\varnothing \left\{ \mathcal{M}_{D_\varphi|^1}^\varnothing \left\{ \dots \left\{ \mathcal{M}_{D_\varphi^{|\mathbb{Q}_\varphi|-1}}^\varnothing \left\{ \mathcal{M}_{D_\varphi}^* \right\} \right\} \right\} \right\}$$

If another value vector D'_φ is subsequently encountered for the first time, the hierarchy of submonitors becomes as follows:

$$\mathcal{M}_{D_\varphi|^0}^\varnothing \left\{ \mathcal{M}_{D_\varphi|^1}^\varnothing \left\{ \dots \left\{ \mathcal{M}_{D_\varphi}^* \right\} \right\}, \mathcal{M}_{D_\varphi'^1}^\varnothing \left\{ \dots \left\{ \mathcal{M}_{D_\varphi'}^* \right\} \right\} \right\}$$

Since the hierarchy is formulated as a recursive set, no duplicate submonitors are allowed. Two submonitors are duplicates, if they represent identical value vectors. If $D_\varphi|^1 = D_\varphi'^1$, the respective monitors are merged. Such merging is explained in detail in Section 4.

3.4 Evaluating $\text{LTL}_4\text{-C}$ Properties

Once the $\text{LTL}_4\text{-C}$ monitor instantiates its submonitors, every submonitor is responsible for updating its truth value. The truth value of LTL_4 submonitors (\mathcal{M}^*) maps to the current state of the submonitor's automaton as described in Definition 6. Quantifier submonitors update their truth value based on the truth values of all child submonitors. The number of child submonitors whose truth value is \top is stored in v_\top (i.e., the \top dimension of vector v) and so on for all truth values in \mathbb{B}_6 . Then, $\text{LTL}_4\text{-C}$ semantics are applied, beginning with function \mathcal{S} (see Equation 6), which in turn relies on the cardinality of function $\mathcal{B}(\varphi, u, D_\varphi|^i, b)$ where b is a truth value. This cardinality is readily provided by the vector v , such that for instance $\mathcal{B}(\varphi, u, D_\varphi|^i, \top) = v_\top$ and so on.

Since each submonitor depends on its child submonitors, updating truth values proceeds outwards, starting at LTL_4 submonitors, then recursively parent submonitors update their truth values until the submonitor $\mathcal{M}_{D_\varphi|^0}^\varnothing$, which is the root submonitor. The truth value of the root submonitor is the truth value of property φ with respect to trace u . This is visualized as the tree shown in Figure 2.

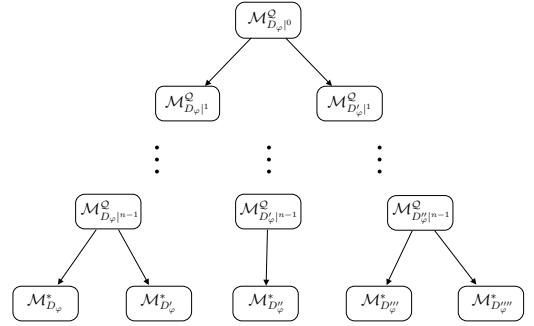


Figure 2: Tree structure of an $\text{LTL}_4\text{-C}$ monitor.

4. PARALLEL ALGORITHM DESIGN

The main challenge in designing a runtime monitor is to ensure that its behavior does not intervene with functional and extra-functional (e.g., timing constraints) behavior of the program under scrutiny. This section presents a parallel algorithm for verification of $\text{LTL}_4\text{-C}$ properties. Our idea is that such a parallel algorithm enables us to offload the monitoring tasks into a different computing unit (e.g., the GPU). The algorithm utilizes the popular *MapReduce* technique to spawn and merge submonitors to determine the final verdict. This section is organized as follows: Subsection 4.1 describes how valuations are extracted from a trace in run time, and Subsection 4.2 describes the steps of the algorithm in detail.

4.1 Valuation Extraction

Valuation extraction refers to obtaining a valuation of quantified variables from the trace. As described in $\text{LTL}_4\text{-C}$ semantics, the predicate $p_i(x_i)$ identifies the subset of the domain of x_i over which the quantifier is applied: namely the subset that exists in the trace. From a theoretical perspective, we check whether the predicate is a member of some trace event, which is a set of predicates. From an implementation perspective, the trace event is a key-value structure, where the key is for instance a string identifying the quantified variable, and the value is the concrete value of the quantified variable in that trace event. Consider the following property:

$$\varphi = \mathbb{A}_{\geq 0.95} s : \text{socket}(s) \Rightarrow (\mathbf{G} \text{recv}(s) \Rightarrow \mathbf{F} \text{respond}(s)) \quad (8)$$

Predicate p in this case is $\text{socket}(s)$, and a trace event should contain a key socket and a value $\in [0, 65535]$ representing the socket file descriptor in the system. Thus, the valuation extraction function $\varepsilon(u_i, K) = D_\varphi$ returns a map where keys are in K , and the value of each key is the value of the quantified variable corresponding to this key. These keys are defined by the user.

4.2 Algorithm Steps

Algorithm 1 presents the pseudocode of the parallel monitoring algorithm. Given an $\text{LTL}_4\text{-C}$ property $\varphi = \mathbb{Q}_\varphi \psi$, the input to the algorithm is the LTL_4 monitor \mathcal{M}^* of LTL_4 property ψ , a finite trace u , the set of quantifiers \mathbb{Q}_φ , and the vector of keys K used to extract valuations. Note that the algorithm supports both online and offline runtime verification. Offline mode is straightforward since the algorithm receives a finite trace that it can evaluate. In the case of online mode, the algorithm maintains data structures that represent the tree structure shown in Figure 2, and repeated invocation of the algorithm updates these data structures incrementally. Thus, a monitoring solution can invoke the algorithm periodically or based on same event in an online fashion, and still receive an evolving verdict.

The entry point to the algorithm is at Line 5 which is

invoked when the monitor receives a trace to process. The algorithm returns a truth value of the property at Line 8. Subsections 4.2.1 – 4.2.4 describe the functional calls between Lines 5 – 8. The MapReduce operations are visible in functions *SortTrace* and *ApplyQuantifiers*, which perform a *map* (\Rightarrow) in Lines 10 and 51 respectively. *ApplyQuantifiers* also performs a *reduction* (\mapsto) in Line 52.

4.2.1 Trace Sorting

As shown in Algorithm 1, the first step in the algorithm is to sort the input trace u (Line 5). The function *SortTrace* performs this functionality as follows:

1. The function performs a parallel map of every trace event to the value vector that it holds using ε (Line 10).
2. The mapped trace is sorted in parallel using the quantifier variable keys (Line 11). For instance, according to Property 8, the key used for sorting will be *socket*, effectively sorting the trace by socket identifier.
3. The sorted trace is then compacted based on valuations, and the function returns a map μ where keys are value vectors and values are the ranges of where these value vectors exist in trace u (Line 12). A range contains the start and end index. This essentially defines the subsequences u^{D_φ} for each property instance $\varphi(D_\varphi)$ (refer to Subsection 2.4).

4.2.2 Monitor Spawning

Monitor spawning is the second step of the algorithm (Line 6). The function *SpawnMonitors* receives a map μ and searches the cached collection of previously encountered value vectors \mathbb{D} for duplicates. If a value vector in μ is new, it creates submonitors and inserts them in the tree of submonitors T (Line 19). The function *AddToTree* attempts to generate $|\mathbb{Q}_\varphi| - 1$ quantifier submonitors \mathcal{M}^φ (Line 26) ensuring there are no duplicate monitors in the tree (Line 27). After all quantifier submonitors are created, *SpawnMonitors* creates an LTL₄ submonitor \mathcal{M}^* and adds it as a child to the leaf quantifier submonitor in the tree representing the value vector (Line 20). This resembles the structure in Figure 2. Creation of submonitors is performed in parallel for all value vectors in trace u .

4.2.3 Distributing the Trace

The next step in the algorithm is to distribute the sorted trace to all LTL₄ submonitors (Line 7). The function *Distribute* instructs every LTL₄ submonitor to process its respective trace by passing the full trace and the range of its respective subsequence, which is provided by the map μ (Line 42). The LTL₄ monitor updates its state according to the trace subsequence and stores its truth value b .

4.2.4 Applying Quantifiers

Applying quantifiers is a recursive process, beginning with the leaf quantifier submonitors and proceeding upwards towards the root of the tree (Line 8). Function *ApplyQuantifiers* operates in the following steps:

1. The function retrieves all quantifier submonitors at the i^{th} level in the tree T (Line 50).
2. In parallel, for each quantifier submonitor, all child submonitor truth values are reduced into a single truth value of that quantifier submonitor (Lines 51-53). This step essentially *reduces* all child truth vectors into a single vector and then applies LTL₄-C semantics to determine the truth value of the current submonitor.

3. The function proceeds recursively calling itself on submonitors that are one level higher. It terminates when the root of the tree is reached, where the truth value is the final verdict of the property with respect to the trace.

Algorithm 1 LTL₄-C monitoring algorithm

```

1: INPUT: An LTL4 monitor  $\mathcal{M}^*$  of LTL property  $\psi$ , a finite trace  $u$ , a
   set of quantifiers  $\mathbb{Q}_\varphi$ , and a vector of keys  $K$  to extract valuations
   of quantified variables.
2: declare  $T = \{\mathcal{M}_{D|0}^\varphi\}$  ▷ Tree of quantifier submonitors
3: declare  $\mathbb{D} = \{\}$  ▷ Value vector set
4: declare  $\mathbb{M}^* = \{\}$  ▷ LTL4 submonitor set
5:  $\mu \leftarrow \text{SORTTRACE}(u)$  ▷ The entry point
6: SPAWNMONITORS( $\mu$ )
7: DISTRIBUTE( $u, \mu$ )
8: return APPLYQUANTIFIERS( $|\mathbb{Q}_\varphi| - 1$ )

```

```

9: function SORTTRACE( $u$ ) ▷ Trace sorting and compaction
10:  $u_i \Rightarrow u'_i := \varepsilon(u_i, K)$  ▷ || map to value vectors
11: PARALLELSORT( $u', K$ )
12:  $\mu(D, r) \leftarrow \text{PARALLELCOMPACT}(u')$ 
13: return  $\mu$ 
14: end function

```

```

15: function SPAWNMONITORS( $\mu$ ) ▷ Monitor spawning
16: for  $D \in \mu$  do in parallel
17:   if  $D \notin \mathbb{D}$  then
18:     ADD( $\mathbb{D}, D$ )
19:      $t \leftarrow \text{ADDTOTREE}(D)$ 
20:      $t.\text{addMonitor}(\text{CREATEMONITOR}(D))$ 
21:   end if
22: end for
23: end function

```

```

24: function ADDTOTREE( $D$ )
25:  $t = T.\text{root}$ 
26: for  $i \in [1, |\mathbb{Q}_\varphi| - 1]$  do
27:   if  $\mathcal{M}_{D|i}^\varphi \notin t.\text{children}$  then
28:      $t.\text{addchild}(\mathcal{M}_{D|i}^\varphi)$ 
29:   end if
30:    $t \leftarrow t.\text{children}[\mathcal{M}_{D|i}^\varphi]$ 
31: end for
32: return  $t$ 
33: end function

```

```

34: function CREATEMONITOR( $D$ ) ▷ Monitor creation
35:  $\mathcal{M}_D^* \leftarrow \text{LAUNCHMONITORTHREAD}(D)$ 
36:  $\mathcal{M}_D^*.D \leftarrow D$ 
37: ADD( $\mathbb{M}^*, \mathcal{M}_D^*$ )
38: return  $\mathcal{M}_D^*$ 
39: end function

```

```

40: function DISTRIBUTE( $u, \mu$ ) ▷ Distribute trace to monitors
41: for  $\mathcal{M}_D^* \in \mathbb{M}^*$  do in parallel
42:   PROCESSBUFFER( $\mathcal{M}_D^*, u, \mu[\mathcal{M}_D^*.D]$ )
43: end for
44: end function

```

```

45: function PROCESSBUFFER( $\mathcal{M}_D^*, u, r$ ) ▷ Process trace subsequence
46: filter include  $u \Rightarrow u' := u[r.\text{start}, r.\text{end}]$  ▷ || filter
47:  $\mathcal{M}_D^*.b \leftarrow \text{UPDATEMONITOR}(\mathcal{M}_D^*, u')$ 
48: end function

```

```

49: function APPLYQUANTIFIERS( $i$ ) ▷ Apply quantifiers
50: for  $t \in T.\text{nodesAtDepth}(i)$  do in parallel
51:    $t.\text{children} \Rightarrow \{s := [v, v', \dots]\}$  ▷ || map to truth vectors
52:    $s \mapsto t.v$  ▷ || reduction to truth vector
53:    $t.b \leftarrow \text{VALUATION}(t)$  ▷ LTL4-C semantics
54: end for
55: if  $i = 0$  then
56:   return  $t.b$ 
57: end if
58: return APPLYQUANTIFIERS( $i - 1$ )
59: end function

```

5. IMPLEMENTATION AND EXPERIMENTAL RESULTS

We have implemented Algorithm 1 for two computing technologies: Multi-core CPUs and GPUs. We applied three

optimizations in our GPU-based implementation: (1) we use *CUDA Thrust API* to implement parallel sort, (2) we using *Zero-Copy Memory* which parallelizes data transfer with kernel operation without caching, and (3) we enforced alignment, which enables coalesced read of trace events into monitor instances. In order to intercept systems calls, we have integrated our algorithm with the Linux **strace** application, which logs all system calls made by a process, including the parameters passed, the return value, the time the call was made, etc. Notice that using **strace** has the benefit of eliminating static analysis for instrumentation. The work in [7, 17, 18] also use **strace** to debug the behavior of applications.

Subsection 5.1 presents the case studies implemented to study the effectiveness of the GPU implementation in online and offline monitoring. Subsection 5.2 discusses the experimental setup, while Subsection 5.3 analyzes the results.

5.1 Case studies

We have conducted the following three case studies:

1. **Ensuring every request on a socket is responded to.** This case study monitors the responsiveness of a web server. Web servers under heavy load may experience some timeouts, which results in requests that are not responded to. This is a factor contributing to the uptime of the server, along with other factors like power failure, or system failure. Thus, we monitor that at least 95% of requests are indeed responded:

$$\mathbb{A}_{\geq 0.95} s : \text{socket}(s) \mathbf{G} \text{ receive}(s) \Rightarrow \mathbf{F} \text{ respond}(s)$$

We utilize the Apache Benchmarking tool to generate different load levels on the Apache Web Server.

2. **Ensuring fairness in utilization of personal cloud storage services.** This case study is based on the work in [10], which discusses how profiling DropBox traffic can identify the bottlenecks and improve the performance. Among the issues detected during this analysis, is a user repeatedly uploading chunks of maximum size to DropBox servers. Thus, it is beneficial for a runtime verification system to ensure that the average chunk size of all clients falls below a predefined maximum threshold, effectively ensuring fairness of service use. The corresponding LTL₄-C property is as follows:

$$\mathbb{A}u : \text{user}(u) \Rightarrow \mathbf{F} (\text{avg_chunksize}(u) \leq \text{maximum})$$

where `avg_chunksize` is a predicate that is based on a variable in the program representing the average chunk size of the current user's session.

3. **Ensuring proxy cache is functioning correctly.** This experiment is based on a study that shows the effectiveness of utilizing proxy cache in decreasing YouTube videos requests in a large university campus [19]. Thus, we monitor that no video is requested externally while existing in the cache:

$$\mathbb{A}v : \text{vid}(v) \Rightarrow \mathbb{E}_{=0} r : \text{req}(r) \Rightarrow (\text{cached}(v) \wedge \text{external}(r))$$

5.2 Experimental Setup

Experiment Hardware and Software. The machine we use to run experiments comprises of a 12-core Intel Xeon E5-1650 CPU, an Nvidia Tesla K20c GPU, and 32GB of RAM, running Ubuntu 12.04.

Experimental Factors. The experiments involve comparing the following factors:

- **Implementation.** We compare three implementations of the LTL₄-C monitoring algorithm:
 - *Single Core CPU.* A CPU implementation running on a single core. The justification for using a single core is to allow the remaining cores to perform the main functionality of the system without causing contention from the monitoring process.
 - *Parallel CPU.* A CPU implementation running on all 12 cores of the system. The implementation uses OpenMP.
 - *GPU.* A parallel GPU-based implementation.
- **Trace size.** We also experiment with different trace sizes to study the scalability of the monitoring solution, increasing exponentially from 16, 384 to 8, 388, 608.

Experimental Metrics. Each experiment results in values for the following metrics:

- **Total execution time.** The total execution time of the monitor.
- **Monitor CPU utilization.** The CPU utilization of the monitor process.

In addition, we measure the following metrics for Case Study 1, since it utilizes an online monitor:

- **Monitored program CPU utilization.** The CPU utilization of the monitored program. This is to demonstrate the impact of monitoring on overall CPU utilization.
- **strace parsing CPU utilization.** The CPU utilization of the **strace** parsing module. This module translates **strace** strings a numerical table.

We perform 20 replicates of each experiment and present error bars of a 95% confidence interval.

5.3 Results

The results of Case Study 1 are shown in Figure 3. As seen in the figure, the GPU implementation scales efficiently with increasing trace size, resulting in the lowest monitoring time of all three implementations. The GPU versus single core CPU speedup ranges from 0.8 to 1.6, increasing with the increasing trace size. When compared to parallel CPU (CPU ||), the speedup ranges from 0.78 to 1.59. This indicates that parallel CPU outperforms GPU for smaller traces (32768), yet does not scale as well as GPU in this case study. This is attributed to the low number of individual objects in the trace, making parallelism less impactful. CPU utilization results in Figure 3 show a common trend with the increase of trace size. When the trace size is small, parallel implementations incur high CPU utilization as opposed to a single core implementation, which could be attributed to the overhead of parallelization relative to the small trace size. On the other hand, GPU shows a stable utilization percentage, with a 78% average utilization. The single core CPU implementation shows a similar trend, yet slightly elevated average utilization (average 86%). The parallel CPU implementation imposes a higher CPU utilization (average 1.15%), since more cores are being used to process the trace. This result indicates that shipping the monitoring workload to GPU consistently provides more time for CPU to execute other processes including the monitored process. The results of Case Study 2 and Case Study 3 in Figures 4 and 5 respectively provide a different perspective. The number of individual objects in these traces are large, making parallelism highly effective. For Case Study 2, the speedup of the GPU implementation over single core CPU ranges from 1.8 to 3.6, and 0.83 to 1.18 over parallel CPU. The average CPU utilization of GPU, single core CPU, and parallel




Figure 3: Results of Case Study 1.




Figure 4: Results of Case Study 2.

CPU is 64%, 82%, and 598% respectively. For Case Study 3, speedup is more significant, with 6.3 average speedup of GPU over single core CPU, and 1.75 over parallel CPU. The average CPU utilization of GPU, single core CPU, and parallel CPU is 73%, 95%, and 680% respectively. Thus, the parallel CPU implementation is showing large speedup similar to the GPU implementation, yet also results in a commensurate CPU utilization percentage, since most cores of the system are fully utilized.

Although the parallel CPU implementation provides reasonable speedup, and the single-core CPU implementation imposes low CPU utilization overhead, the GPU implementation manages to achieve both simultaneously.

6. RELATED WORK

Runtime verification of parametric properties has been studied by Rosu et al [11, 12, 15]. In this line of work, it is possible to build a runtime monitor parameterized by objects in a Java program. The work by Chen and Rosu [8] presents a method of monitoring parametric properties in which a trace is divided into slices, such that each monitor operates on its slice. This resembles our method of identifying trace subsequences and how they are processed by submonitors. However, parametric monitoring does not provide a formalization of applying existential and numerically constrained quantifiers over objects.

Bauer et al. [5] present a formalization of a variant of first order logic combined with LTL. This work is related to our work in that it instantiates monitors at run time according to valuations, and defines quantification over a finite subset of the quantified domain, normally with that subset being defined by the trace. Our work extends this notion with numerical constraints over quantifiers, as well as a parallel algorithm for monitoring such properties.

The work by Leucker et al. presents a generic approach for monitoring modulo theories [9]. This work provides a more expressive specification language. Our work enforces a canonical syntax which is not required in [9], resulting in more expressiveness. However, the monitoring solution provided requires SMT solving at run time. This may induce




Figure 5: Results of Case Study 3.

substantial overhead as opposed to the lightweight parallel algorithm presented in this paper, especially since it is designed to allow offloading the workload on GPU. SMT solving also runs the risk of undecidability, which is not clear whether it is accounted for or not. LTL_4-C is based on six-valued semantics, extending LTL_4 by two truth values: \top_c and \perp_c . These truth values are added to support quantifiers and their numerical constraints. This six-valued semantics provides a more accurate assessment of the satisfaction of the property based on finite traces as opposed to the three-valued semantics in [9]. Finally, although LTL_4-C does not support the expressiveness of full first-order logic, numerical constraints add a flavor of second-order logic increasing its expressiveness in the domain of properties where some percentage or count of satisfied instances needs to be enforced.

The work in [1] presents a method of using MapReduce to evaluate LTL properties. The algorithm is capable of processing arbitrary fragments of the trace in parallel. Similarly, the work in [2] presents a MapReduce method for offline verification of LTL properties with first-order quantifiers. Our work uses a similar approach in leveraging MapReduce, yet also adds the expressiveness of counting semantics with numerical constraints. Also, our approach supports both offline and online monitoring by introducing six-valued semantics, which are capable of reasoning about the satisfaction of a partial trace. This is unclear in [2], since there is no evidence of supporting online monitoring.

Finally, the work in [6] presents two parallel algorithms for verification of propositional LTL specifications at run time. These algorithms are implemented in the tool RiTHM [16]. This paper enhances the framework in [6, 16] by introducing a significantly more expressive formal specification language along with a parallel runtime verification system.

7. CONCLUSION

In this paper, we proposed a specification language (LTL_4-C) for runtime verification of properties of types of objects in software and networked systems. Our language is an extension of LTL that adds counting semantics with numerical constraints. The six truth values of the semantics of LTL_4-C allows system designers to obtain informative verdicts about the status of system properties at run time. We also introduced an efficient and effective parallel algorithm with two implementations on multi-core CPU and GPU technologies. The results of our experiments on three real-world case studies show that runtime monitoring using GPU provides us with the best throughput and CPU utilization, resulting in minimal intervention in the normal operation of the system under inspection.

For future work, we are planning to design a framework for monitoring LTL_4-C properties in distributed systems and cloud services. Another direction is to extend LTL_4-C such that it allows non-canonical strings of quantifiers. Finally, we are currently integrating LTL_4-C in our tool RiTHM [16].

8. REFERENCES

- [1] Benjamin Barre, Mathieu Klein, Maxime Soucy-Boivin, Pierre-Antoine Ollivier, and Sylvain Hallé. Mapreduce for parallel trace validation of ltl properties. In *Runtime Verification*, pages 184–198. Springer, 2013.
- [2] David Basin, Germano Caronni, Sarah Ereth, Matúš Harvan, Felix Klaedtke, and Heiko Mantel. Scalable offline monitoring. In *CADE*, pages 1–16. Springer, 2013.
- [3] A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL semantics for runtime verification. *Journal of Logic and Computation*, 20(3):651–674, 2010.
- [4] A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL Semantics for Runtime Verification. *Journal of Logic and Computation*, 20(3):651–674, 2010.
- [5] Andreas Bauer, Jan-Christoph Küster, and Gil Vegliach. From propositional to first-order monitoring. In *Runtime Verification*, pages 59–75. Springer Berlin Heidelberg, 2013.
- [6] S. Berkovich, B. Bonakdarpour, and S. Fischmeister. GPU-based runtime verification. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1025–1036, 2013.
- [7] Maximiliano Caceres. Syscall proxying-simulating remote execution. *Core Security Technologies*, 2002.
- [8] Feng Chen and Grigore Roşu. Parametric trace slicing and monitoring. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 246–261. Springer, 2009.
- [9] Normann Decker, Martin Leucker, and Daniel Thoma. Monitoring modulo theories.
- [10] Idilio Drago, Marco Mellia, Maurizio M Munafo, Anna Sperotto, Ramin Sadre, and Aiko Pras. Inside dropbox: understanding personal cloud storage services. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 481–494. ACM, 2012.
- [11] Soha Hussein, Patrick Meredith, and Grigore Roşu. Security-policy monitoring and enforcement with javamop. In *Proceedings of the 7th Workshop on Programming Languages and Analysis for Security, PLAS ’12*, pages 3:1–3:11, New York, NY, USA, 2012. ACM.
- [12] Dongyun Jin, P.O. Meredith, Choonghwan Lee, and G. Rosu. Javamop: Efficient parametric runtime monitoring framework. In *Software Engineering (ICSE), 2012 34th International Conference on*, pages 1427–1430, June 2012.
- [13] Leonid Libkin. *Elements of finite model theory*. springer, 2004.
- [14] Z. Manna and A. Pnueli. *Temporal verification of reactive systems - safety*. Springer, 1995.
- [15] Patrick Meredith and Grigore Rosu. Efficient parametric runtime verification with deterministic string rewriting. In *Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on*, pages 70–80. IEEE, 2013.
- [16] S. Navabpour, Y. Joshi, C. W. W. Wu, S. Berkovich, R. Medhat, B. Bonakdarpour, and S. Fischmeister. Rithm: a tool for enabling time-triggered runtime verification for c programs. In *ACM Symposium on the Foundations of Software Engineering (FSE)*, pages 603–606, 2013.
- [17] Daniel Ramsbrock, Robin Berthier, and Michel Cukier. Profiling attacker behavior following ssh compromises. In *Dependable Systems and Networks, 2007. DSN’07. 37th Annual IEEE/IFIP International Conference on*, pages 119–124. IEEE, 2007.
- [18] Feng Wang, Qin Xin, Bo Hong, Scott A Brandt, Ethan L Miller, Darrell DE Long, and Tyce T McLarty. File system workload analysis for large scale scientific computing applications. In *Proceedings of the 21st IEEE/12th NASA Goddard Conference on Mass Storage Systems and Technologies*, pages 139–152, 2004.
- [19] Michael Zink, Kyoungwon Suh, Yu Gu, and Jim Kurose. Watch global, cache local: Youtube network traffic at a campus network: measurements and implications. In *Electronic Imaging 2008*, pages 681805–681805. International Society for Optics and Photonics, 2008.